arXiv:1408.6621v1 [cs.HC] 28 Aug 2014

Tuning the Diversity of Open-Ended Responses From the Crowd

Walter S. Lasecki [*1], Christopher M. Homan [†2], and Jeffrey P. Bigham [‡3]

[1]University of Rochester, Computer Science, Rochester, NY 14627, wlasecki@cs.rochester.edu
[2]Rochester Institute of Technology, Computer Science, Rochester, NY 14623, cmh@cs.rit.edu
[3]Carnegie Mellon University, HCI and LTI Institutes, Pittsburgh, PA 15213, jbigham@cs.cmu.edu

August 29, 2014

Abstract

Crowdsourcing can solve problems that current fully automated systems cannot. Its effectiveness depends on the reliability, accuracy, and speed of the crowd workers that drive it. These objectives are frequently at odds with one another. For instance, how much time should workers be given to discover and propose new solutions versus deliberate over those currently proposed? How do we determine if discovering a new answer is appropriate at all? And how do we manage workers who lack the expertise or attention needed to provide useful input to a given task? We present a mechanism that uses distinct payoffs for three possible worker actions—*propose*, *vote*, or *abstain*—to provide workers with the necessary incentives to guarantee an effective (or even optimal) balance between searching for new answers, assessing those currently available, and, when they have insufficient expertise or insight for the task at hand, abstaining. We provide a novel game theoretic analysis for this mechanism and test it experimentally on an image-labeling problem and show that it allows a system to reliably control the balance between discovering new answers and converging to existing ones.

Introduction

The rapidly increasing popularity of crowdsourcing is due in large part to its demonstrated capability to perform tasks that lie far beyond the reach of current fully au-

[*]Supported in part by a Microsoft Research Graduate Fellowship.
[†]Supported in part by NSF Award #SBE-1111016.
[‡]Supported in part by NSF Award #IIS-1149709.

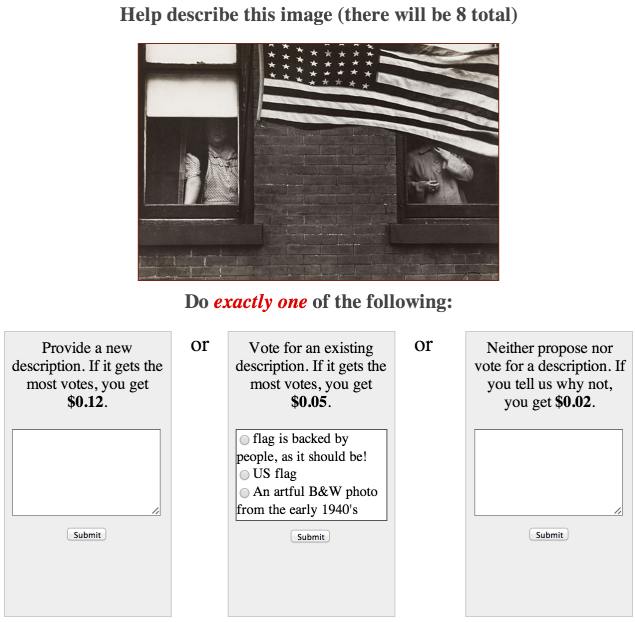
Help describe this image (there will be 8 total)
Do exactly one of the following:
Provide a new description. If it gets the most votes, you get $0.12.
Submit
or
Vote for an existing description. If it gets the most votes, you get $0.05.
flag is backed by people, as it should be!
US flag
An artful B&W photo from the early 1940's
Submit
or
Neither propose nor vote for a description. If you tell us why not, you get $0.02.
Submit

Figure 1: An example of the sort of problems that the propose-vote-abstain can solve. Workers are asked to label a series of images. They are shown potential payoffs, and asked to choose between voting, proposing, or abstaining from contributing an answer.

tomated systems, and to marshal human resources more efficiently than offline approaches can. Prior work has shown that crowdsourcing can complete tasks such as editing documents [BLM+10], labeling images [vAD04], discovering the folding structure of protein [CKT+10], discovering celestial bodies [LSS+08], and evaluating the quality of products and services [AKM12]. Recent advances have even led to real time services [LMW+11, LWN+13, BBMK11, AIP11, AIT11], such assistance technology that answers visual questions [BJJ+10, LTZ+13] or help find accessible paths [HLF13] for blind users, and captions audio in any setting for deaf and hard of hearing users [LMS+12].

Across this wide range of applications, crowdsourcing relies on the mystery of human intelligence and the highly distributed nature of the crowd to gain its advantages. This leads to two basic design problems: how to motivate the crowd to contribute good information and how to integrate this information into a coherent solution. Eliciting and integrating multiple workers responses has become a key attribute of many crowd-powered systems – for instance, prior work has investigated how to coordinate and aggregate the input of multiple workers to determine the best edits to make in a document [BLM+10], combine real-time interface control input [LMW+11], and determine the best response to provide in a conversation [LWN+13].

The relationship between data production and filtering is complex. In general, more workers contributing data means more filtering is necessary, which increases the overall cost of the task, in terms of (human or otherwise) computing resources. On the other hand, if too little effort is spent on discovery then the results may be unsatisfying. Additionally, if as here we assume that humans perform both tasks then we need to respect the competencies and preferences of these workers and route them to the particular tasks that they are most fit to perform.

In this paper we explore the use of game-theoretic principals to achieve the right balance between data production and filtering. We present the *propose-vote-decline* mechanism. Each crowd worker is given a choice among *proposing* an answer, *voting* among the answers proposed so far, or *abstaining*, i.e., doing nothing. When a stopping condition is reached, the mechanism returns the answer with the most votes. Workers are paid a base amount, with bonuses if they propose or vote for the winning answer.

Much of the game theoretic research on crowdsourcing falls into one of three categories. The first studies crowdsourced contests, in which participants compete in some sort of creative task, such as discovering the folding structure of protein [CKT+10] (where workers compete for high scores) or design a new collaborative filtering algorithm for Netflix [BK07] (where workers competed for a cash prize). This research typically views crowdsourcing as an all-pay auction [DV09]. Studies have compared the performance (in terms of market efficiency) of winner-take-all versus distributed reward structures [BLL11, MS01, MS06, AS09, CHS12]. In most of this work, the amount of time in which a crowd worker is engaged can be very long, the value function for the contributions is public and objective, and which worker(s) are rewarded is determined by a principal actor. The major strategic problem the workers face is thus to determine how much effort they should expend. In this work, we are also interested in determining the best contribution(s).

In contrast, however, we are interested in situations in which evaluation is subjective in nature and there may be a great deal of uncertainty about what the "best" con-

tribution is. We assume that the wisdom of the crowd is the best arbiter of quality—an assumption that is valid in many settings. We also assume that workers are engaged for short periods of time and paid a small amount for their contributions. These assumptions hold true in many crowdsourcing applications. Yet under these assumptions it is not realistic for workers to choose strategically the amount of effort they will expend on their contributions, as the range of effort levels available to choose from is quite narrow. Rather, workers can choose strategically whether to propose a new contribution, endorse (by voting for) an existing contribution, or abstain altogether, and our goal is to design mechanisms that provide the best answer available on demand and quickly converge to an optimal answer.

The second major category of game theory applied to crowdsourcing deals with recommender systems. Like the present paper, much this work [RV97, Pre04, MRZ05, JF05, JF07, WP12], assumes that rating is a subjective process and use voting to perform evaluation. Unlike our work, much of this research assumes that the alternatives to evaluate are predetermined. A notable exception is Gao et al. [GBKG12] who, like us, consider jointly the proposal and voting processes. However, unlike us, they consider both processes as distinct with distinct workers assigned to distinct roles as voters or authors, and where the voters evaluate the authors, not their contributions. By contrast, we allow the workers to determine the roles they wish to play and rather consider the effect different incentives have on the roles they chose.

Nearly all of the previously-mentioned work assumes that the workers are synchronized and that they play their moves simultaneously. By contrast, we assume that workers are asynchronous and play sequentially. This makes the analysis somewhat more complex, however it allows to study real-time concerns.

The third major stream of game-theoretic research considered control-theoretic mechanisms, primarily in settings where workers are engaging over relatively long periods of time, as in citizen science. Kamar et al. use such partially-observeable Markov decision processes to study a crowdsourcing problem in which the workers procede in a series of rounds, and describe an efficient algorithm that, in their case, learns an optimal policy [KHH12].

Social choice theory has also been used in the context of crowdsourcing, say, to power collaborative filtering. Parkes and Procaccia study social choice functions [PP13] embedded in Markov decision processes. These systems consider game-theoretic behavior from the requester's rather than the worker's perspective. We, on the other hand, are concerned with the strategies of the workers.

Hemaspaandra et al. [HHR14] consider online voting where the agents vote sequentially and each agent may know the votes of the preceding agents, but not the succeeding ones. The alternatives to choose from, however, are know in advances to all agents. They are primary concerned with manipulation in the setting. In our setting, the alternatives available change dynamically over time. Also, we do not directly study manipulation in this paper. We discuss both of this issues in the limitations section.

Like most social choice research, both works above assume that all alternatives are known to all the beforehand. By contrast, the dynamic discovery of the alternatives (proposals) is a central goal of our system, and so we cannot assume that the alternatives are known beforehand.

The Propose-Vote-Abstain Mechanism

Each propose-vote-abstain *round* begins with a *request*—for information or an answer to a specific question—by the *requester*. Workers are then recruited into the system.

As soon as each worker joins, he or she is presented with the request and all the contributions proposed so far by the previous workers. The new worker is then offered the following choices: *propose* a new contribution, for a payoff of π if that answer eventually receives the most votes, *vote* for one of the contribution that has already been proposed, for a payoff of ν if that contribution eventually receives the most votes, or abstain for an unconditional payoff of α. If at termination more than one contribution has the most votes, then one of them is selected uniformly at random as the winner.

Crucially, no worker knows the number of votes each existing contribution has received, nor how many workers have already played, nor how many workers total have been recruited. This helps curb the *majority effect*, though obviously a strategic player would vote for the candidate believed to be the one most preferred by the other voters.

The process continues until some stopping condition is reached. This could be after a fixed number of workers respond, or one (or more) contribution receives a certain minimum number of votes.

In this paper, we assume that the payoff structure (π, ν, α) is fixed per round (i.e., the same for all workers) and only consider single round settings, i.e., each worker makes exactly one move (vote, propose, abstain). Multiround settings are also natural to consider, especially in the context of highly interactive, deliberative group problem solving, or in realtime streaming situations when the mechanism must continuously emit solutions (or receive new requests based on previous solutions). However, single round settings are common enough to use them in this first study on the mechanism.

It is also reasonable to consider variations in which (π, ν, α) vary dynamically within a single round. Likewise, the requirement that only one alternative be chosen or that plurality is the social choice function can be generalized. Such considerations are beyond the scope of this paper.

Modeling Worker Behavior

From the theoretical perspective we view the propose-vote-abstain mechanism as a single round, sequential-play, extensive-form game. Doing so helps us to reason about how different incentive structures may lead to different outcomes. Here we provide a number of baseline results, and give a few glimpses into how the same framework can be used to model real-world behavior. Our assumptions are inspired by a scenario that is arguably not likely to occur in practice, but which here provides very useful baseline.

Catastrophic Freeloading

We call this scenario *Catastrophic Freeloading*. It describes a situation in which no worker is willing to put any effort into proposing or voting, but is willing to randomly vote or propose if doing so yields a higher expected payoff than abstaining. We describe

it in terms of a belief model in which each worker is completely oblivious to the choices available and, furthermore, believe that all other workers are similarly obviously.

This scenario is important to consider because it describes the amount of data produced by a completely ignorant and unmotivated (but strategically clever) workforce and estimates the cost (in worker payments) to the system for that data.

In particular, we make the following assumptions.

1. The game has a single turn with a indeterminate (and unknown to the workers) number of players.
2. Workers play sequentially, one at a time.
3. The game does not terminate unless there is at least one answer proposed and one vote.
4. We assume that all players have equal confidence in any alternative being better than another, including those the players themselves may propose.
5. The only information the workers know about the system are the candidate answers and the intial request for proposals.
6. The tasks required of the worker are simple and require a fixed time and energy commitment.

Assumptions (4)–(5) mean effectively that the only information we have to distinguish one state (i.e., set of candidate answers) from another is the number of alternatives m_t proposed so far, where t is the current time. (4) lets us assume that the votes cast at any time are uniformly distributed at random from among all alternatives available at that time. That is, *we assume here that voting is effectively random selection*. Essentially, this means that all workers are completely ignorant about the alternatives available to them, and they believe that all the other workers share this level of ignorance. Thus, the most effective way to choose an alternative is by blind guessing. It ignores shared biases, basic and expert knowledge, irrational behavior, and many other factors that might lead to more complex behavior.

These catastrophic freeloading assumptions are obviously very strong, so much so that any notion of "preference" over the set of proposals is all but erased. We take it to be a baseline condition, against which assumptions accounting for worker preferences can be measured. This is, admittedly, a small step, but a necessary one, for how to model worker preferences and—as is crucial in game theoretic considerations—*beliefs about other workers' preferences*—is not a straightforward problem, nor one that can be easily generalized across a potentially broad spectrum of potential applications.

An alternative view of these assumptions is they reflect the actions of a worker that is completely disengaged in the process, i.e., "phoning it in." It could be seen as an extreme case where all workers are freeloaders who want to maximize their expected reward while spending a minimum amount of time.

Regarding (6), in prior, auction-based research [MS01] workers had complete information about the game and were allowed to choose strategically the amount of effort to expend on the problem. Under those assumptions the purpose of the mechanism was

to provide inventives for the voters to provide the right amount of effort. In microtask settings, such as the one on which our experiment is based, workers tend to be engaged for short amounts of time and are not expected to bring much effort to any one transaction. Thus, from a design perspective it is much more important to focus on the *beliefs* and *knowledge* the workers can bring to the task right now, as opposed to the amount of effort they could spend on more open-ended tasks.

With these assumptions in mind, Theorem 1 tells us quite a bit about the baseline behavior of the crowd, from a mechanism design perspective. In particular, it considers, given a

Theorem 1. *Given some target number of candidate proposals $m \in \mathbb{N}$, any cost structure (π, ν, α) that satisfies $(m+1) \cdot \alpha > \nu > m \cdot \alpha$ and $(m+1) \cdot \nu > \pi > m \cdot \nu$ generates a dominant strategy in which the first m workers propose and the remaining workers vote.*

Theorem 1 follows from the propositions below. It tells us that, when workers are maximally uncertain about the choices at hand then (1) the abstention payoff provides an incentive for limiting the number of proposals, *though not the number of votes*; and (2) the workers monotonically propose first and vote later, never returning again to propose.

Before we present these propositions, let us first provide some useful notation and observations. Let m_t denote the number of candidate proposals present for the tth worker to choose from. Note that under the above assumptions this is effectively all that the worker observes about the candidate proposals, as he or she is completely oblivious to the content of each proposal. Thus, all we need to know about the state in order to model worker behavior is m_t for each worker t. Thus, for any other worker t', if $m_t = m'_t$ then t and t' have the same dominant strategy. In particular, if some dominant strategy results in $m_{t+1} = m_t$ then all remaining workers will repeat this action for the rest of the game. This is the case whenever the dominant strategy at m_t is to vote or to abstain. If the dominant strategy is to propose, then $m_{t+1} = m_t + 1$. This simple observation is key to understanding the propositions.

Let $S^*(m_t)$ denote the dominant strategy for m_t.

Proposition 2 provides basic results about worker strategies, in particular sheds light on the role that abstention plays.

Proposition 2. *Under assumptions (4)–(5):*

1. *If $\alpha \geq \min\{\pi, \nu\}$ then abstaining is a dominant strategy for all players.*
2. *If $\min\{\pi, \nu\}/(m_t + 1) < \alpha$ then for worker t abstaining has a greater expected payoff than proposing.*
3. *If $\nu/m_t < \alpha$ then for worker t abstaining has a greater expected payoff than voting.*

Proof. (1) is obvious.

(2)—(3) follow from assumption (4), in that having uniform confidence in each candidate–and assuming that all workers feel similarly—means that any voter is equally

likely to chose any candidate currently available, so that the expected payoff for proposing or voting is upper bounded by dividing the base reward π or ν among the number of proposals available to vote on after the worker moves, which is $m_t + 1$ if worker t proposes and m_t if worker t votes. For (2), note additionally that whenever proposing would create so many candidates as to give voting a lower expected payoff than abstaining, proposing has a lower expected payoff than abstaining. □

Next, we study conditions that lead to proposal and voting.

Proposition 3. *If $\nu > \pi > \alpha$ then $S^*(1) = \text{vote}$.*

Proposition 3 is obvious.

While the earlier propositions provide important boundary results, Proposition 4 provides insight on the system state as the game unfolds. The proposition works backwards, but finding a state in which workers vote. Since voting does not change the state, the system will remain in this state until termination. We then use structural induction to work backwards from this voting state to characterize the rest of the state space.

Proposition 4. *If $\pi > \nu > \alpha$ then:*

1. *If $m_t > \nu/\alpha$, then $S^*(m_t) = \text{abstain}$.*
2. *If $m_t + 1 > \nu/\alpha > m_t > 0$ then $S^*(m_t) = \text{vote}$.*
3. *For all $t, w, m_t, m_w \in \mathbb{N}$ such that $w > t$, $m_w > m_t > 0$, $S^*(m_w) = \text{vote}$ and, for all $u : m_w > m_u > m_t$, it holds that $S^*(m_u) = \text{propose}$:*
   (a) *If $m_t > \nu \cdot m_w/\pi$ then $S^*(m_t) = \text{propose}$.*
   (b) *If $m_t < \nu \cdot m_w/\pi$ then $S^*(m_t) = \text{vote}$.*

   Let m_0^ν denote $\lfloor \nu/\alpha \rfloor$ and, for all $i \geq 0$, let $m_{i+1}^\nu = \lfloor \pi \cdot m_i^\nu/\nu \rfloor$.
4. *For all $i \geq 0$ if $m_i^\nu > 0$ then $S^*(m_i^\nu) =$ vote.*
5. *For all m_t such that $m_t \leq \nu/\alpha$ and $m_t \notin \{m_i^\nu\}$, $S^*(m_t) =$ propose.*

Proof. Item (1) follows from Proposition 2.3.

For item (2), if $m_t + 1 > \nu/\alpha > m_t$ then voting has a higher expected payoff than abstaining, as long as no one else proposes. But proposing would make the expected payoff for voting less than for abstaining. Consequently, the expected payoff for proposing is zero at this point (assuming all subsequent workers act to optimize their expected payoffs). Thus voting has the highest expected payoff.

For item (3), if $m_t > \nu \cdot m_w/\pi$ then π/m_w, the expected payoff for proposing is greater than ν/m_t the expected payoff for voting, which is greater than α (otherwise $S^*(m_w) \neq$ vote). If $m_t < \nu \cdot m_w/\pi$ then by a similar argument the expected payoff for voting is greater than proposing.

Items (4) and (5) follows by applying item (3) inductively to each m_i^ν.

□

Items (4) and (5) show that even in the relatively small state space we consider here that the space of dominant strategies is fairly rich, varying between proposing and voting for different values of m_t. However, little of it is reachable in practice, as the state does not change one workers start voter.

Finally Proposition 5 provides initial conditions.

Proposition 5. *If* $\alpha < \min\{\pi, \nu\}$ *then* $S^*(0) =$ *propose.*

Proof. If worker 0 proposes then its expected payoff is greater than abstaining or voting as long as there are never more than $\min\{\nu, \pi\}/\alpha$ proposals. But this will never happen because of Proposition 2.2. □

With these results in mind, we can now prove Theorem 1.

Proof (of Theorem 1). For a given target number of proposals m and cost structure (π, ν, α) that satisfies $(m+1) \cdot \alpha > \nu > m \cdot \alpha$ and $(m+1) \cdot \nu > \pi > m \cdot \nu$, note that the only $i \in \mathbb{N}$ for which $m_i^\nu > 0$ is $i = 0$. Thus, by Propositions 4.4–refit:whenever-propose, $S^*(m) = \text{vote}$ and for all m' between 0 and m (exclusive), $S^*(m') = \text{propose}$. By Proposition 5, $S^*(0) = \text{propose}$. □

What conclusions should a system designer draw from these results? Certainly, in order to making voting or proposing incentive compatible actions, it is obvious that voting must pay more than abstaining, and proposing must pay more than voting. It is also reasonable to assume that the more one pays for each action, the more effort workers will spend to perform them. These results show that there is a another side to the matter: that raising incentives for proposing and voting can also increase the amount of bad data gathered. Of course, the catastrophic freeloader scenario we study here may not apply often in practice. Even when there are freeloaders, they may not dominate the system, and even when they do, it is not clear that the freeloaders would themselves be aware of their dominant situation and so might not adopt the strategies seen here (that is, they might behave differently if they believed that some of the other workers are putting effort and attention into their work).

The results also suggest that conscientious, informed workers will act when freeloaders may not, and so, to avoid freeloading, the best approach is to make voting pay more than abstention, but less than twice as much, and make proposing pay twice as much as voting, but less than twice as much. Under such a cost structure, Theorem 1 says that under the catastrophic freeloading modeling, one freeloading might propose (and any number of them could vote as long as there is only one proposal), but any remaining proposals and votes would come from workers who more sincerely believe in their choices.

Experiments

We recruited 111 Mechanical Turk workers and asked them to view a set of 5 images (presented in random order) and either propose, vote for, or abstain from contributing to the image's description. A round consists of a single image with a fixed (i.e., the same for all workers in the round) payoff structure (i.e., fixed prices) for proposing,

π	ν	workers	proposals	votes	abstains
\$0.20	\$0.04	20	39	60	1
\$0.12	\$0.05	33	41	117	7
\$0.08	\$0.08	20	13	86	1
\$0.05	\$0.12	10	13	36	0
\$0.04	\$0.20	18	13	76	3

Table 1: Basic statistics from the experiments. The payment for abstaining is fixed at \$0.02. Each HIT asks each worker about the same five images, shown in Figure 2, for proposal and voting payoffs π and ν, respectively.

voting, and abstaining. All five rounds assigned to each worker had the same price structure There were in total five different payoff structures, leading to 25 rounds total.

Table 1 shows the five payoff structures and how the workers and their actions distribute over them.

Figure 2 shows the images to be labeled. In order to not bias the experiments towards specific domains, the images were chosen to be vague.

Figure 3 shows the results of our tests on the number of proposals generated. As the payment for voting becomes large relative to the proposal payment, the number of total answers generated by the system significantly decreases from an average of 7.8 responses to 1 response for all 5 pictures we saw ($p < .0001$). The decreasing trend was linear with $R^2 = 0.802$. Note that there is a disproportional drop at the break-even point when payment is equal for both options. This is consistent with what we expect because voting requires less effort than generating a response, so there is a slight bias in its favor.

While all of the images eventually converged to a single response as the vote payment increased, the number of responses generated in the opposing case (where the proposal reward is high and workers are incentivized to generate several answers) varies from 5 to 10 responses each. This likely depends on how subjective the image is and how many answer could be considered plausible with high confidence. This trend is seen throughout the results as each response set trends towards a single response. This suggests that the content of the task does play a role in worker trends, but in the convergent limit this can be overridden by financial incentives.

Overvoting

One problem we did not expect to encounter to such a degree was overvoting (mainly because we assumed that workers would expect for us to check for this). It is possible for workers who choose to vote to do so more than once to do so simply by hitting the "back" button on their browser. It is not clear whether this happens accidentally or intentionally, however, Table 2 shows the total number of over- versus actual- votes in the rounds of each price structure. The amount of overvoting in the second and fifth price structures in particular seems too high to be solely accidental.

IMG1

IMG2

IMG3

IMG4

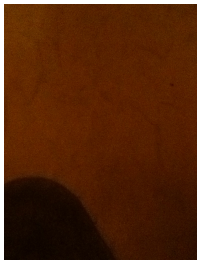

IMG5

Figure 2: The images used in the experiments.

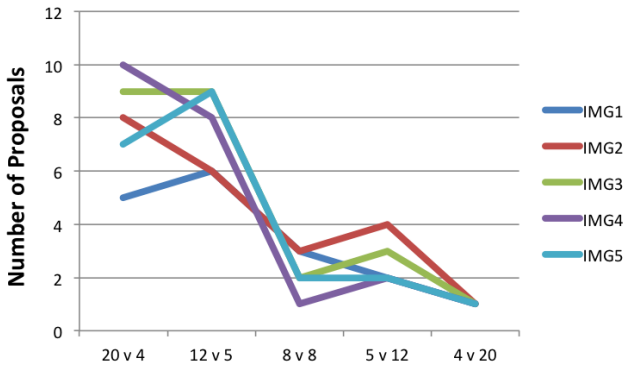

Figure 3: Results from varying relative payments for voting and proposing responses to an image description task.

Theoretical Predictions

How did the workers in our experiment fare against the theoretical predictions? The theory predicts that whenever the payoffs for proposing and voting are greater than the payoff for abstaining then the first worker will propose and at least one worker will vote. This was consistent with our experiments.

It also predicts that when the payoff for abstaining is greater than either voting or proposing then no one will vote or propose. We used such a payoff structure in one of our training rounds, and in none of those cases did a worker propose or vote.

Additionally, the theory predicts that there would be no abstentions. Table 1 shows that, in fact, few workers ever abstained.

Moving beyond these basic results, the theory predicts that, under the catastrophic freeloader, there are two different values that upper bound the number of proposals: namely π/α and ν/α, and $\min\{\pi/\nu, 1\}$, where π, ν, and α are the payoffs for, respectively, proposing, voting and abstaining. Table 3 shows the number of rounds where the number of proposals exceeds at least one of these bounds.

These results, though admittedly of a preliminary nature, might support the hypothesis that π/α is an effective upper bound on the proposals. In terms of our worker's beliefs (and under standard bounded rationality assumptions), this could suggest that, when it comes to proposing new contributions, workers have little confidence in their proposals.

The results also show that the remaining bounds—ν/α, and $\min\{\pi/\nu, 1\}$—having do with how with the expected payoff to future voters and thus about what the current worker thinks of the competence of future workers to vote for his or her contribution, if the current worker proposes, are frequently violated. This would seem to suggest that workers, even though they are not confident about their ideas, are rather more confident about the ability of future voters to make the right choice.

π	ν	number of voting workers	number of actual votes
\$0.20	\$0.04	60	69
\$0.12	\$0.05	117	147
\$0.08	\$0.08	86	90
\$0.05	\$0.12	36	36
\$0.04	\$0.20	76	124
	total	375	464

Table 2: The number of votes registered for each round in each payoff structure compared to the number of workers who choose to vote. Overvoting occurs when workers reload a screen and can be accidental or intentional. Lines two five seem to show evidence of cheating. The payment for abstaining is fixed at \$0.02.

		# rounds where # proposals exceeds		
π	ν	π/α	$\min\{\pi/\nu, 1\}$	ν/α
20	4	0	5	0
12	5	2	5	5
8	8	0	0	4
5	12	0	0	2
4	20	1	0	4
	total	3	10	15
	winners	0	4	4

Table 3: How does actual worker behavior compare to the theoretical predictions under the zero knowledge assumption? This table looks at three values that upper bound the number of proposals under those predictions. The values are based on the payoffs for proposing, π, voting, ν, and abstaining α. Note that the maximum number of rounds for each price structure is five.

Another takeaway is that the number of times a winning contribution proposed after any one of the three boundaries is violated is very small (eight total). The theory suggests that a rational worker would violate these boundaries only if the worker was confident in having an answer. Our results might suggest that such a worker's confidence is often misplaced.

Regarding monoticity, we looked at the order in which votes and proposals appear. Rather than all proposals appear first, followed by all votes, with no interleaving, votes and proposals were in face interspersed. Consequently, later proposals were eligible for fewer votes. Figure 4 show the percentage of voting workers who did so after the last contribution was proposed.

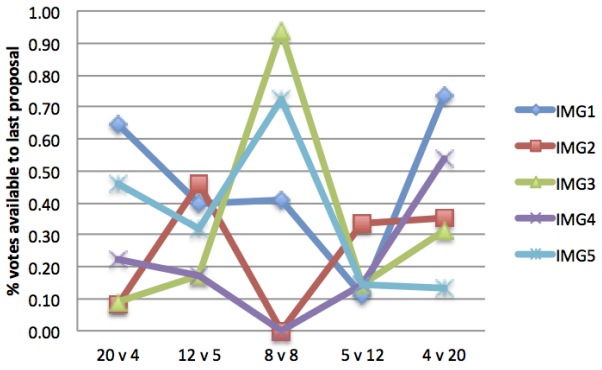

Figure 4: Percentage of votes available to the last contribution. In our experiments we observed that—contrary to prediction—workers do not all propose first and then vote after all proposals appear; proposals and votes are interleaved. Thus, proposals that appear earlier are generally eligible for more votes than those that appear later.

Discussion

Among our incoming assumptions was that the time commitment and financial benefits to microtasking are so small that workers do not make decisions about the amount of effort to put into a task. The possible existence of cheating among voters suggests that worker effort, even at this low-stakes level, must still be considered. We are curious about the motivations of the cheaters. Do they think that cheating gives them a leg up on their competitors, or do they assume that everyone else is doing and so they need to follow suit just to keep up with their peers?

From a theoretical perspective, voting systems in which the alternatives to choose from and votes themselves are added asynchronously, in sequence, presents a number of interesting challenges. For instance, in general there is a bias towards the alternatives that appear earlier, since they have a long time frame within which to gather votes. There are several ways to mitigate this bias, e.g., by fixing the interval of time between when a contribution is proposed and when it can last be voted on, so that all contributions have the same amount of time to garner votes. One can also use importance sampling or impose halting conditions that guarantee each proposal gets a fair chance of being selected. How to best do so depends on the application at hand and is a matter of further study.

Limitations

There are a number of limitations to our approach. The first is that game theory generally imposes a number of simplifying assumptions on the behavior of individuals. This is usually necessary to make the theory tractable, but often leads to behavior that is not observed in practice. Our view hear is that it can nonetheless go a long way in explaining observed behavior if it is regarded as presenting an ideal case (or null hypothesis)

and we reflect on *how* (or if) observed behavior deviates from it.

Another limitation is that we do not consider the possibility that workers could collude, say, to propose an answer and then all vote for it. Clearly, the incentive mechanism does not account for this. Voting is well known to be vulnerable to manipulation under ideal conditions, however, how vulnerable they are in practice is a matter of ongoing research. Another way around this problem is to start with a baseline assumption about how large such a manipulating coalition can be and then make the worker pool large enough to negate any effect the manipulators might have. Finally, one can user worker feedback systems to punish those who manipulate through collusion.

Another limitation is that we are taking the "best" answer to be that chosen by the majority of the workers. Though this can be effective in certain circumstances—such as putting general labels on images—but not as well when expert advice is needed. For instance, if a worker happened to know know that IMG2 was a by the artist Robert Frank, would he or she propose this information? Or would he or she hold back and propose something simpler (or not propose at all), expecting this fact to be lost on the average worker?

While we might expect the latter behavior to dominate, we at least once observed the former behavior. In one round involving IMG5, the proposal that received the most votes was a later one that read, "Metal gate blocking a stone staircase that leads up to an ancient aqueduct." Perhaps the specificity of the answer conveyed enough authority to convince voters to support it. (Incidently, the image is of the Colosseum of Rome, so the proposal is not too far off the mark.)

The rest of the proposals from this round are illuminating to read, as they convey a variety of strategies in proposing answers. They read, in order of appearance, "gate", "Protect the place", "Staircase", "Fenced steps", (the winning proposal), "steps to prison", "This is a metal grate installed over the staircase to one of Rome's deteriorating ruins" (this proposal is slightly less specific than the winning one, and also less erroneous), "Ruins are ruins, why want to peep in? Hence the grill-grate!!" (which has a less authoritative tone than the winning answer, and almost seems to be chiding some of the other answers for their specificity), "METAL GATE." Perhaps the difference between the winning and losing proposals here is that the level of detail (even if false) and language used in the winning post suggests a higher level of authority in the image than the losing proposals.

How to account theoretically for more realistic worker beliefs is a subject of future research, but we sketch some ideas here (we also explore this question experimentally in the next section). We can for instance assume a belief model where the worker at time t has beliefs above the relative likelihood of each proposal succeeding so far. We can represent those beliefs as real numbers $w_{m_0}, \ldots, w_{m_t} \in [0, 1]$, where t_m is the number of proposals at time t m_t. Let $w' \in [0, 1]$ represent the worker's confidence in his or her own choice. Finally let P_t be a distribution over $[0, 1]$ from which future worker's submissions are drawn. The likelihood that the if the worker votes he or she chooses the contribution proposed at time $t' < t$ is $m'_t / \sum_{0 \leq i \leq w_{m_t}} w_i$. Under the zero knowledge assumption, each of these weights is identical.

Conclusion

In this paper we discussed and formalized a model for eliciting varying levels of diversity in contributions from crowd workers. While the individual components of this contribution are built on common building blocks. The tradeoff in response set size that we have demonstrated has not been demonstrated before to our knowledge. Our results suggest future work in exploring why and in what circumstances workers will self-select out of a task, and provides a basis for exploring mechanism-based self-filtering of responses by workers.

References

[AIP11] Josh Attenberg, Panagiotis G. Ipeirotis, and Foster J. Provost. Beat the machine: Challenging workers to find the unknown unknowns. In *Human Computation*, volume WS-11-11 of *AAAI Workshops*. AAAI, 2011.

[AIT11] Sinan Aral, Panos Ipeirotis, and Sean Taylor. Content and context: Identifying the impact of qualitative information on consumer choice. In Dennis F. Galletta and Ting-Peng Liang, editors, *ICIS*. Association for Information Systems, 2011.

[AKM12] Omar Alonso, Gabriella Kazai, and Stefano Mizzaro. *Crowdsourcing for Search Engine Evaluation*. Springer, 2012.

[AS09] Nikolay Archak and Arun Sundararajan. Optimal design of crowdsourcing contests. In *ICIS*, page 200, 2009.

[BBMK11] M. S. Bernstein, J. Brandt, R. C. Miller, and D. R. Karger. Crowds in Two Seconds: Enabling Realtime Crowd-powered Interfaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2011)*, pages 33–42, Santa Barbara, CA, 2011.

[BJJ+10] J. P. Bigham, C. Jayant, H. Ji, G. Little, A. Miller, R. C. Miller, R. Miller, A. Tatarowicz, White S. White, B., and T. Yeh. VizWiz: nearly real-time answers to visual questions. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2010)*, pages 333–342, New York, NY, 2010.

[BK07] Robert M Bell and Yehuda Koren. Lessons from the netflix prize challenge. *ACM SIGKDD Explorations Newsletter*, 9(2):75–79, 2007.

[BLL11] Kevin J Boudreau, Nicola Lacetera, and Karim R Lakhani. Incentives and problem uncertainty in innovation contests: An empirical analysis. *Management Science*, 57(5):843–863, 2011.

[BLM+10] Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. Soylent: a word processor with a crowd inside. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, UIST '10, pages 313–322, New York, NY, USA, 2010. ACM.

[CHS12] Shuchi Chawla, Jason D. Hartline, and Balasubramanian Sivan. Optimal crowdsourcing contests. In Yuval Rabani, editor, *SODA*, pages 856–868. SIAM, 2012.

[CKT+10] Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker, Zoran Popovic, and Foldit Players. Predicting protein structures with a multiplayer online game. *Nature*, 466(7307):756–760, 2010.

[DV09] Dominic DiPalantino and Milan Vojnovic. Crowdsourcing and all-pay auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 119–128. ACM, 2009.

[GBKG12] Xi Alice Gao, Yoram Bachrach, Peter Key, and Thore Graepel. Quality expectation-variance tradeoffs in crowdsourcing contests. In *AAAI*, 2012.

[HHR14] Edith Hemaspaandra, Lane A Hemaspaandra, and Jörg Rothe. The complexity of online manipulation of sequential elections. *Journal of Computer and System Sciences*, 80(4):697–710, 2014.

[HLF13] Kotaro Hara, Vicki Le, and Jon Froehlich. Combining crowdsourcing and google street view to identify street-level accessibility problems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 631–640, New York, NY, USA, 2013. ACM.

[JF05] Radu Jurca and Boi Faltings. Enforcing truthful strategies in incentive compatible reputation mechanisms. In *Internet and Network Economics*, pages 268–277. Springer, 2005.

[JF07] Radu Jurca and Boi Faltings. Robust incentive-compatible feedback payments. In *Agent-Mediated Electronic Commerce. Automated Negotiation and Strategy Design for Electronic Markets*, pages 204–218. Springer, 2007.

[KHH12] Ece Kamar, Severin Hacker, and Eric Horvitz. Combining human and machine intelligence in large-scale crowdsourcing. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '12, pages 467–474, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

[LMS+12] W. S. Lasecki, C. D. Miller, A. Sadilek, A. AbuMoussa, and J.P. Bigham. Real-time captioning by groups of non-experts. In *In Submission*, 2012. http://hci.cs.rochester.edu/pubs/pdfs/legion-scribe.pdf.

[LMW+11] W. S. Lasecki, K. I. Murray, S. White, R. C. Miller, and J. P. Bigham. Real-time Crowd Control of Existing Interfaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2011)*, pages 23–32, Santa Barbara, CA, 2011.

[LSS+08] Chris J Lintott, Kevin Schawinski, Anže Slosar, Kate Land, Steven Bamford, Daniel Thomas, M Jordan Raddick, Robert C Nichol, Alex Szalay, Dan Andreescu, et al. Galaxy zoo: morphologies derived from visual inspection of galaxies from the sloan digital sky survey. *Monthly Notices of the Royal Astronomical Society*, 389(3):1179–1189, 2008.

[LTZ+13] Walter S. Lasecki, Phyo Thiha, Yu Zhong, Erin Brady, and Jeffrey P. Bigham. Answering visual questions with conversational crowd assistants. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '13, pages 18:1–18:8, New York, NY, USA, 2013. ACM.

[LWN+13] Walter S Lasecki, Rachel Wesley, Jeffrey Nichols, Anand Kulkarni, James F Allen, and Jeffrey P Bigham. Chorus: a crowd-powered conversational assistant. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, pages 151–162. ACM, 2013.

[MRZ05] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

[MS01] Benny Moldovanu and Aner Sela. The optimal allocation of prizes in contests. *American Economic Review*, pages 542–558, 2001.

[MS06] Benny Moldovanu and Aner Sela. Contest architecture. *Journal of Economic Theory*, 126(1):70–96, 2006.

[PP13] David C Parkes and Ariel D Procaccia. Dynamic social choice with evolving preferences. pages 767–773, 2013.

[Pre04] Dražen Prelec. A bayesian truth serum for subjective data. *Science*, 306(5695):462–466, 2004.

[RV97] Paul Resnick and Hal R Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, 1997.

[vAD04] Luis von Ahn and Laura Dabbish. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '04, pages 319–326, New York, NY, USA, 2004. ACM.

[WP12] Jens Witkowski and David C Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 964–981. ACM, 2012.